

Golestani 3

capacity, however, that corresponds to the link capacity divided by a "utilization factor." See the equation for  $\Delta$  in FIG. 1 of Afek et al.

The aim is to reduce the rates of sessions (sharing the link) that are above  $\Delta$ , toward  $\Delta$ , and to increase the rates of sessions (sharing the link) that are below  $\Delta$ , also towards  $\Delta$  (see col. 6, lines 29-34). Once a switch computes the  $\Delta$  of a link, it modifies the Maximum Allowed Cell Rate (MACR) in accordance with a specified equation (e.g., the equation at col. 7, lines 67 of Afek et al).

Flow control is implemented in an ATM network with a resource management cell (RM cell) that loops around the virtual circuit of each session. The RM cell contains several fields, and one of them is the Explicit Rate (ER) field. The flow control is implemented by controlling the value of the ER field as follows: if the ER field of a backward traveling RM cell exceeds the current value of MACR for the link through which the RM cell is about to exit a switch, then the ER field is set equal to the MACR value of the link. This way, when the RM cell returns to the session's source, the lowest MACR value encountered in the path is revealed to the source and, with this information available, the source can perform calculations that lead to a control of the source's rate of transmission. Afek et al do NOT disclose exactly how the source controls its rate of transmission, i.e., what those calculations are.

In a TCP network, the source periodically polls the routers on the path to the destination. In each polling, the sources fetches the MACR values of the links of those routers, and adjusts its window according to the minimum MACR (co. 10, lines 46-51). Again, Afek et al do NOT disclose the calculations that lead to the control of its window size. In another implementation that is suited for TCP networks where the header is modified to include a field that contains the current rate of the source, the routers set an appropriate bit in such a field of those packets with a rate that is above MACR, and the reaction of the source to the receipt of such packets (with a set bit) is as if those packet were dropped, decreasing its rate of transmission.

There doesn't seem to be any explicit teaching in Afek et al as to where the method of their invention is practiced, in the sense of where the computations shown in FIG. 1 are made. However, one might presume that the computations are performed in each switch, since there is no suggestion that information necessary for the computation of  $\Delta$ , MACR,

Golestani 3

etc. – which is available in the switches – is transmitted to any other place, and since no approach is disclosed as to how the computed MACR values may be communicated back to the switches to affect the RM cell's ER field.

Regarding the rejected claims, the following addresses the claims in order of their rejection, starting with claim 1, which is the parent claim of all of the rejected claims.

First, the method described by Afek et al is link-centric, and is most likely executed by, and within, each switch/router of the network, as explained above. The computed MACR value, for example, is a link congestion measure.

In contradistinction, claim 1 defines a method that is session-centric (is carried out by a session). That alone patentably distinguishes claim 1 from the Afek et al reference.

Second, the first step of the method defined in claim 1 states:

evaluating a session congestion measure that is related to congestion information on links of said network which carry incoming traffic to a receiving end of said session.

The Examiner asserts that Afek et al describe this step in col. 11, lines 6-10, but applicant contends that Afek et al do not evaluate a session congestion measure at all, and the text in col. 11, lines 6-10 does NOT support this assertion by the Examiner. Rather, the cited text merely states:

3. Selective set of EPCI (Explicit Forward Congestion Indication) bit:  
The route sets the EPCI bit in packets whose indicates rate is above MACR. The reaction of the source to the receipt of a packet with a set EPCI bit is analogous to that when receiving a source quench message

but that, however, merely relates to a mechanism by which a router signals to the source the value of MACR that the router computed. As indicated above, *the MACR value is a measure that relates to a link, not to a session.*

In contradistinction, the first step specifies evaluating a session congestion measure from "information on links" (such as MACR values, or value), and such evaluating is simply not addressed by Afek et al. Hence, the first step of claim 1 also directs the conclusion that claim 1 is neither anticipated nor rendered obvious by the Afek et al reference.

Third, the second step of claim 1 relates to the evaluation of a session incremental reward function. The Examiner equates the session incremental reward function to the MACR, but, as indicated above, the MACR is basically a measure of link congestion, and

Golestani 3

is certainly not a session measure. Hence, it is respectfully submitted that Afek et al neither teach nor suggest the second step of claim 1.

Fourth, the third step of the method defined in claim 1 states:

evaluating a new rate of said incoming traffic that moves said rate of said incoming traffic in a direction that minimizes a global network cost function which combines cost functions assigned to said sessions and congestion cost functions assigned to said links.

The Examiner asserts that Afek et al evaluate a new rate, citing col. 8, line 25, through col. 9, line 19. Applicant respectfully disagrees. There is no discussion in the cited text that pertains to the actual computing of a new rate. Rather, the cited text basically discusses the computing of weights  $\alpha_{inc}$  and  $\alpha_{dec}$ , *which are used in the computation of the MACR*. The MACR is NOT the new rate of packet/cell transmissions.

Moreover, it is respectfully submitted that the third step is not anticipated by the reference on three additional counts: (1) neither in the cited text nor in any other portion of the reference is there any teaching of the actual evaluation of a new rate, (2) neither in the cited text nor in any other portion of the reference is there any teaching of a "global network cost function," and (3) neither in the cited text nor in any other portion of the reference is there any teaching of a cost function "which combines cost functions assigned to said sessions and congestion cost functions assigned to said links."

In short, the above provides numerous reasons for holding (based on each reason individually) that claim 1 is neither anticipated nor is rendered obvious by Afek et al.

As for claim 5, it specifies that the new rate is computed by incrementing the old rate, where the "incrementing is determined based on said session incremental reward function and said session congestion measure." As indicated above, there is no teaching in the reference as to how the new rate is computed. That, by itself, is sufficient to hold that Afek et al do not teach, or suggest, the method of claim 5.

Additionally, what IS known is that the only information that is made available to the unit that does compute the new rate (impliedly, the source of packets) is the MACR. Since claim 5 requires both an "incremental reward function" and a "session congestion measure" to compute a new rate, since the only measure available to the computing unit is the MACR, and since the Examiner equates the MACR with the "incremental reward function," it remains that Afek et al teach the use of a single factor, rather than two and,

Golestani 3

therefore, there is no evidence, or suggestion, that the computations of a new rate includes the use of a "session congestion measure." Indeed employing this measure but not mentioning it would be a violation of the "best mode" requirement. Accordingly, claim 5 contains a limitation that provides support on two grounds for holding that claim 5 is neither anticipated nor rendered obvious by the Afek et al reference (in addition to the claim 1 grounds).

The Examiner asserts that the text in col. 10, lines 7-45 teaches that "the change in the rate of flow of traffic is updated in increments of ^," and applicant recognizes that through a typographical error the symbol "^" probably stands for something else. However, regardless for what it stands, applicant respectfully submits that the Examiner is in error. The cited text deals with the computation of MACR, and the Fast\_MACR. Neither of them is a transmission rate, or a session transmission rate, computation. Further, while it is true that the MACR and the Fast\_MACR are modified incrementally, it does not necessarily follow that new rates are changed incrementally, even if they are based on variables that change incrementally. It remains, therefore, that Afek et al simply do not teach HOW the new rates are computed.

As for claim 6, it adds detail to the claim 1 definition by specifying that:

...said step of evaluating a new rate is carried out at a receiving end of said session, and said method further comprises a step of communicating information to a sending end of said session, to change said rate of said incoming traffic towards said new rate.

The Examiner asserts that col. 10, lines 46-54 discloses this step because, as the Examiner sees it, the evaluation of the new rate of flow is carried out in a router (which the Examiner equates to the "receiving end"), and the source (which is the Examiner equates to the "sending end") polls the router to receive new rate and adjust its window. Respectfully, applicant disagrees.

First, no artisan would assert that a router of a network is a "receiving end" of any communication, and certainly it is quite clear that the term "receiving end" of "a session," in the context of claim 6 and in light of applicant's disclosure, cannot be interpreted to be a router of the network.

Second, the router (even if erroneously viewed as a receiving end) does not compute a new rate but, as indicated above, computes an updated MACR, which is NOT a

Golestani 3

new rate -- and even that is a surmise, since Afek et al do not explicitly state what unit computes the MACR.

Thus, it is respectfully submitted that, in addition to the limitations of the parent claim, claim 6 includes limitations that, independently, direct a conclusion that claim 6 is neither anticipated nor rendered obvious by Afek et al.

Regarding claim 7, it is respectfully submitted that the arguments made in connection with claim 6 apply with equal vigor to claim 7.

As for claim 8, it specifies that the "new rate developed is an incremental change arrived at through an additive factor." The Examiner asserts that this limitation is taught in col. 10, lines 9-10 of the reference. Applicant respectfully disagrees.

Lines 9-20 of col. 9 state:

a second weighted average of  $\Delta$  is computed, in the same way as MACR, except that, order...

It is quite clear in the reference that  $\Delta$  is unused link capacity. It is NOT a rate, or the "new rate." Further, the cited language does not teach or suggest anything about the new rate. Certainly, it does not teach or suggest that the new rate "is an incremental change," and as indicated above, the fact that a variable which is employed in calculating the new rate changes incrementally does not dictate that the new rate would change incrementally. Lastly, it does not teach or suggest that the new rate is an incremental change "arrived at through an additive factor."

In short, applicant respectfully submits that the text cited by the Examiner does not support the assertion that Afek et al teach the limitation of claim 8. Claim 8 is believed, therefore, to be neither anticipated nor rendered obvious by Afek et al.

Regarding claim 25, it defines multipath routing where the packets of a particular session, in the course of their flow from their source point to their destination point, are split at some point into two flows that traverse different links, and rejoin at the destinations, or at some switch prior to the destination. Claim 25 specifies this condition by defining that packets of the traffic of a session are split into two sets, that one set traverses a subset of links, and that the other set ("at least some other of said packets") traverse a different subset of links.

In rejecting the claim, the Examiner cites col. 10, lines 63-65, which states:

Golestani 3

...implementation is used: the router indicates the ER on backward packets, and the source station adjusts the window size according to the received ER.

The relevance of this passage to claim 23 is not clear, and the Examiner's comment regarding this passage ("incoming tracks comprises packets") is not understood. The Examiner also cites col. 12, lines 7-13 -- which is claim 3 of Afek et al. This claim merely states that the two links through which a forward cell traverses (one incoming to a switch and one outgoing from a switch) require a fixed transmission time to transmit a cell (packet), and that a previously defined (and above-mentioned) variable,  $\tau$ , is a power of 2 times the fixed transmission time. Neither of these passages has anything to do with splitting a flow into two streams, where packets of one stream traverse a different subset of links than packets the other stream. Hence, it is respectfully submitted that claim 23 is neither anticipated nor rendered obvious by the reference.

Nevertheless, in order to make the multipath definition clearer, claim 23 is amended herein. It is believed that claim 23, as amended, also is not anticipated or rendered obvious by the reference.

It is believed that claims 25-27 draw their patentability from claim 1.

As for claim 32, its rejection must be an error, since it depends on allowed claim 30.

Regarding claim 34, the Examiner asserts that Afek et al teach that the "step of evaluating said session congestion measure equates said session congestion measure to the value of said at least one congestion field of a received probe packet," citing col. 11, lines 37-42. Applicant respectfully disagrees on a number of counts.

First, the cited text (extended slightly to complete the sentence) merely states:

This approach enhances fairness by penalizing sessions whose rate exceeds MACR. An advantages of this implementation is that both the delay and the throughput of "well-behaved" sessions, whose rate is less than MACR, is not affected by the "misbehaved" sessions, whose rate is larger than MACR.

This passage does not support any conclusion regarding session congestion measure, or the equating of session congestion measures to anything.

Second, while it is true that, according to one embodiment, the value of the ER field in RM cells is equated to the value of MACR, it must be remembered that the MACR

Golestani 3

is not a session congestion measure. As indicated above, the MACR is a link congestion measure. In contradistinction, claim 34 specifies a session congestion measure evaluation. Further, it is noted that claim 34 addresses the computations performed in the source, and that Afek et al do not explicitly teach anything that is computed/evaluated in the source (other than though general statements such as "adjusts the window size according to the received ER," col. 10, lines 64-65). Accordingly, it is believed that claim 34 is neither anticipated by nor rendered obvious by Afek et al.

Claims 2-4, 21, and 35 were rejected under 35 USC 103 as being unpatentable over Afek et al, and further in view of Mitra et al, US Patent 6,331,986. Applicant respectfully traverses.

Claim 2 specifies that the

session incremental reward function is the negative of a derivative, with respect to rate of said incoming traffic, of said one of said cost functions assigned to said session.

This requires cost functions as defined in claim 1, an incremental reward function, and the attribute that the incremental reward function is the negative of the derivative of the cost functions.

In claim 1, the Examiner equates the term "incremental reward function" with the NARC of the reference (page 3, third paragraph of the Examiner's remarks). Independent of the fact that the NACR is not an incremental reward function, it is clear that the definition of NACR is not the negative of a derivative of any cost function. Nowhere in Afek et al is there a suggestion any of the other terms in the definition of NACR ( $\Delta$ ,  $1$ ,  $\alpha_{inc}$  and  $\alpha_{dec}$ ) are an integral of a cost function.

The Examiner admits this deficiency in the Afek et al reference, but asserts that the Mitra reference teaches determining traffic rate, and that Mitra et al teach a revenue sensitive to link capacity, pointing to equation 15 in col. 17 of Mitra et al. This equation, indeed, is a revenue sensitivity equation with respect to link capacity and NOT to traffic rate. Therefore, incorporating the teachings of Mitra et al in the Afek et al scheme would still not yields the limitation defined in claim 2. Moreover, it is not at all clear that equation 15 of Mitra et al is actually evaluated in the course of the Mitra et al method (flow charts of FIG. 9 through FIG. 13) and, therefore, it is not clear that the equation 15 even qualifies as a cost function, or a "session incremental reward function." Further the

Golestani 3

calculations of Mitra et al are carried out in an effort to determine a best route through which information can flow from the source to the destination. It is not an effort to control a source's rate. In other words, the Mitra et al work is directed to a different aspect of network control – that of properly allocating traffic that is to be placed on the network, rather than controlling the flow of existing traffic to avoid congestion. Therefore, it is respectfully submitted claim 2 is not obvious in view of the combination of Afek et al and Mitra et al.

Regarding claim 3, a similar argument applies. The Examiner points to equation 3 of Mitra et al, in col. 15, but it is not at all clear that this equation is actually evaluated in the course of the method carried out by Mitra et al, as described in flow charts of FIG. 9 through FIG. 13. Further, the sensitivity measure of equation 3 does not correspond to a “of a sum of congestion cost functions assigned to links employed by said session.” Accordingly, combining the teachings of Afek et al with Mitra et al still does not yield the method of claim 3 and, therefore, applicant respectfully submits that claim 3 is not obvious in view of the Afek et al – Mitra et al combination of references.

With reference to claim 4, the Examiner points to col. 6, lines 38-39 of Mitra et al, and to col. 12, lines 24-27 of Mitra et al. The col. 6 citation (extended to line 42), merely states that the  $c_{ij}$  parameters are referred to as implied costs, and that they

reflect the effective loss of revenue occasioned when the carrying of calls of a given service calls on a given link reduces the remaining capacity and thus leads to an incremental increase in the blocking of future coffered calls.

That simply says the Mitra et al employ “implied cost” parameters. The col. 12 citation teaches that there is a lower bound on link capacity  $C_L$ , i.e.,  $n_L$  bar, an upper bound on link capacity  $C_U$ , i.e.,  $n_U$  double bar, and a parameter that corresponds to link-capacity increments, and thresholds  $H_L$  and  $W_L$  that are used for testing the convergence of the network revenue  $W$ . Neither of the citations, however, teach or suggest assigning a very large cost function to links that carry a load that is greater than a predetermined value, which is what claim 4 specifies. Accordingly, it is believed that claim 4 is not obvious in view of the Afek et al – Mitra et al combination.

As for claim 35, the Examiner points to col. 7 lines 30-32 of Mitra et al, and to col. 11, lines 53-55 of Mitra et al. The text in col. 7 states that complexity of calculating the



Golestani 3

link loss probabilities can be reduced to  $O(1)$  in box 90 of the prior art FIG. 6, but that does not teach, or suggest, that the session congestion measure is based on probability of packet loss experienced at the receiving end. The text in col. 11 states that a solver of the optimal routing problem applies a stochastic model, but to estimate the link-loss probabilities and their partial derivatives with respect to offered traffic, it treats small capacity links by exact techniques, and large-capacity links by asymptotic techniques. That, of course, is not a teaching, or a suggestion that the session congestion measure is based on probability of packet loss experienced at the receiving end. Hence, it is respectfully submitted that claim 35 is not obvious in view of the Afek et al – Mitra et al combination.

In addition to the above remarks, applicant respectfully points out that the Mitra reference does not supply any of the failings of the Afek et al reference discussed above in connection with claim 1 and, therefore, none of the claims rejected under 35 USC 103 in view of Afek et al and Mitra et al are obvious in view of this combination of references.

Claims 12, 14 and 15 were rejected under 35 USC 103 as being unpatentable over Afek et al in view of Szentesi, US Patent 5,844,886. Applicant respectfully traverses. Szentesi describes an arrangement where a central network controller 40 controls blocking of calls within switches of the network. FIG. 3, which the Examiner cites, relates to revenue versus traffic for the **entire network**. It is not an incremental reward function for a **session**, which is the subject matter of claim 12.

As for claims 14 and 15, the Examiner cites FIG. 9 of Szentesi, but this FIG. presents merely simulation results. It does not represent a link cost function that is used in the control of traffic rates by a source.

It is noted, moreover, that the Szentesi reference does not supply that which is missing in the Afek et al reference, as discussed above in connection with claim 1. Hence, it is respectfully submitted that claims 12 and 14-15 are not obvious in view of the Afek et al – Szentesi combination of references.

Golestani 3

In light of the above amendments and remarks, applicants respectfully submit that all of the Examiner's objections and rejections have been overcome. Reconsideration and allowances of all of the outstanding claims are respectfully solicited.

Respectfully,  
Jamaloddin S Golestani

Dated: \_\_\_\_\_

By Henry T. Brendzel  
Henry T. Brendzel  
Reg. No. 26,844  
Phone (973) 467-2025  
Fax (973) 467-6589  
email [brendzel@comcast.net](mailto:brendzel@comcast.net)

Golestani 3

**Appendix - Amendment showing changes made****IN THE SPECIFICATION:**

*Please replace the paragraph beginning at page 4 line 1 with the following:*

It should be noted that for any fixed window size, as the network becomes congested and the round trip delay increases, the transmission rate is concomitantly reduced. This reaction takes place within one round trip delay; i.e. it is[s] a short-term operation. Thus, the window scheme provides a form of dynamic congestion control even if the window size is not adjusted according to network conditions. If modifying the window size in response to quasi-static network conditions is permitted, then the window scheme combines dynamic and quasi-static congestion control. In such an arrangement, the window size can be set to the product of the *medium-term* average rate,  $r_s$ , and the medium term average round trip delay,  $\tau_s$ ; i.e.,  $w_s = r_s \cdot \tau_s$ .

*Please replace the paragraph beginning at page 10 line 11 with the following:*

Realization of this fact suggests that  $\nu_s$  can be used as a priority assignment to sessions. Sessions with larger  $\nu_s$  are cut less severely in response to network congestion. Correspondingly, a larger  $\nu_s$  makes sessions less sensitive to the number of hops they must traverse in the network. It should be mentioned, perhaps, that any advantage gotten from setting  $\nu_s$  at some level is only relative. If all sessions are assigned a large  $\nu_s$ , the congestion measures  $\gamma_s$  will increase until every body is cut back to the proper usage level.

**IN THE CLAIMS:**

*Please amend the claims below to read as follows:*

23. (Amended) The method of claim 1 where said incoming traffic comprises packets where [some] a subset of said packets traverse [one] a first subset of links of said network, [and at least some others of said] remaining packets of said incoming traffic traverse a [different] second subset of links, and said first subset and second subset are mutually exclusive.

Golestani 3

34. (Amended) The method of claim [1] 25 where said step of evaluating said session congestion measure equates said session congestion measure to the value of said at least one congestion field of a received probe packet.